


Grid2003: Using Grids for Science

Paul Avery
University of Florida
avery@phys.ufl.edu

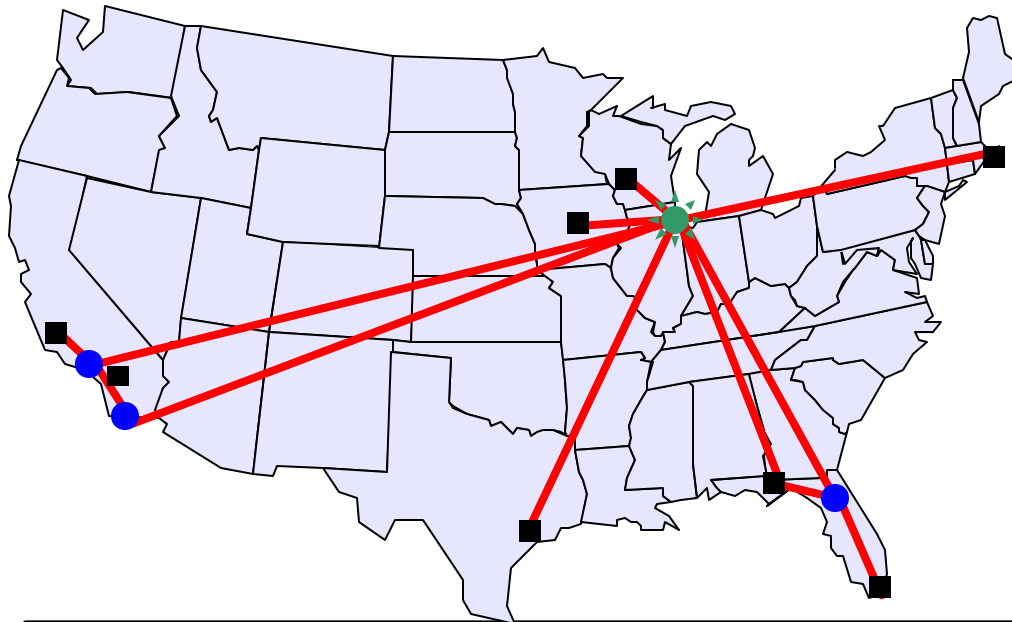


**HEPAP Meeting
Washington, DC
February 10, 2004**



The Grid Concept

- ✂ Grid: Geographically distributed computing resources configured for coordinated use
 - ✂ Fabric: Physical resources & networks provide raw capability
 - ✂ Ownership: Resources *controlled* by owners and *shared* w/ others
 - ✂ Middleware: Software ties it all together: tools, services, etc.
- ✂ Goal: Transparent resource sharing



US-CMS
"Virtual Organization"



Science Drivers for U.S. HEP Grids

✍ LHC experiments
✍ 100s of Petabytes

✍ Current HENP experiments
✍ ~1 Petabyte (1000 TB)

✍ LIGO
✍ 100s of Terabytes

✍ Sloan Digital Sky Survey
✍ 10s of Terabytes

Future Grid resources

✍ Massive CPU (PetaOps)
✍ Large distributed datasets (>100PB)
✍ Global communities (1000s)

2009

2007

2005

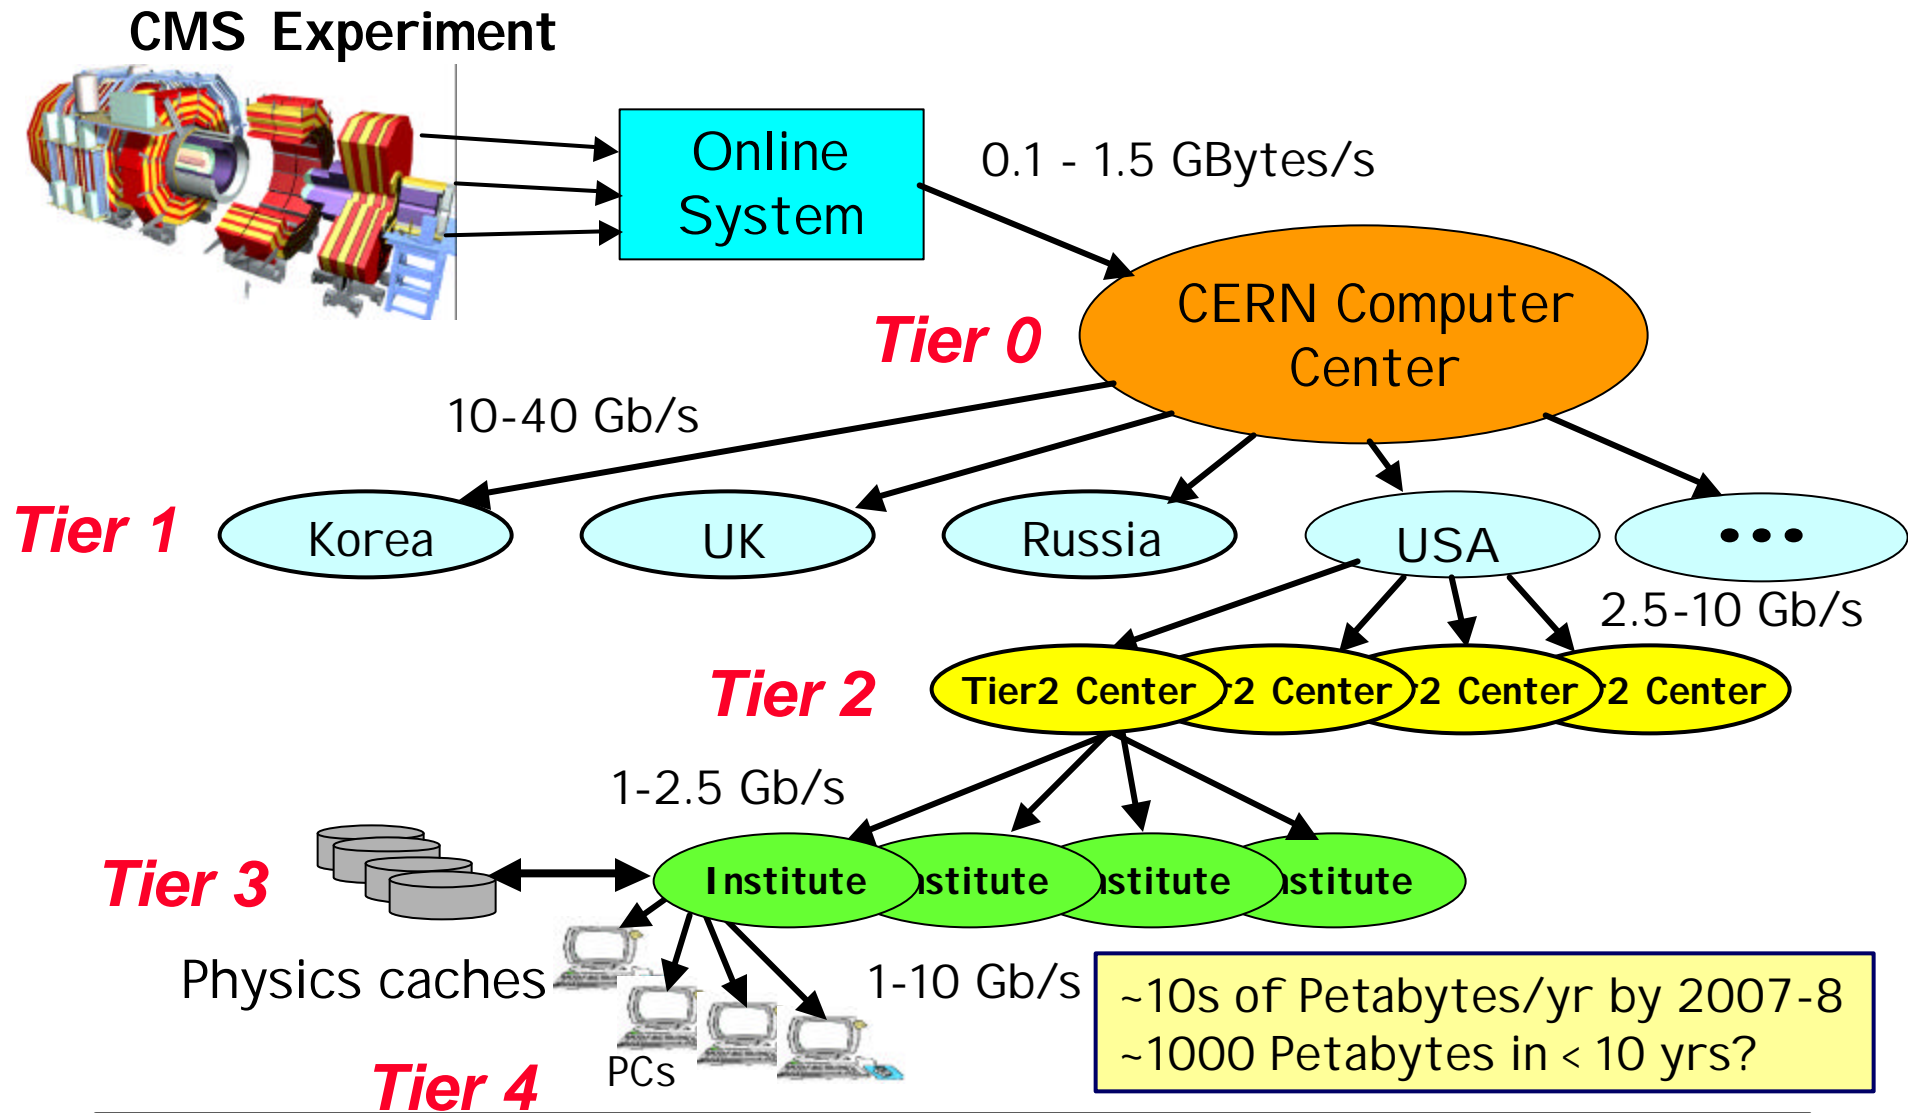
2003

2001

Data growth
Community growth



Global LHC Data Grid Hierarchy





Global Context of Data Grid Projects

Main HEP projects

U.S. Projects

- ✍ GriPhyN (NSF)
- ✍ iVDGL (NSF)
- ✍ Particle Physics Data Grid (DOE)
- ✍ PACs and TeraGrid (NSF)
- ✍ DOE Science Grid (DOE)
- ✍ NEESgrid (NSF)
- ✍ NSF Middleware Initiative (NSF)

EU, Asia projects

- ✍ European Data Grid (EU)
- ✍ EDG-related national Projects
- ✍ DataTAG (EU)
- ✍ LHC Computing Grid (CERN)
- ✍ EGEE (EU)
- ✍ CrossGrid (EU)
- ✍ GridLab (EU)
- ✍ Japanese, Korea Projects

- ✍ **Not exclusively HEP**
- ✍ **But most driven/led by HEP (with CS)**
- ✍ **Many 10s x \$M brought into the field**



"Trillium": U.S. Physics Grid Projects

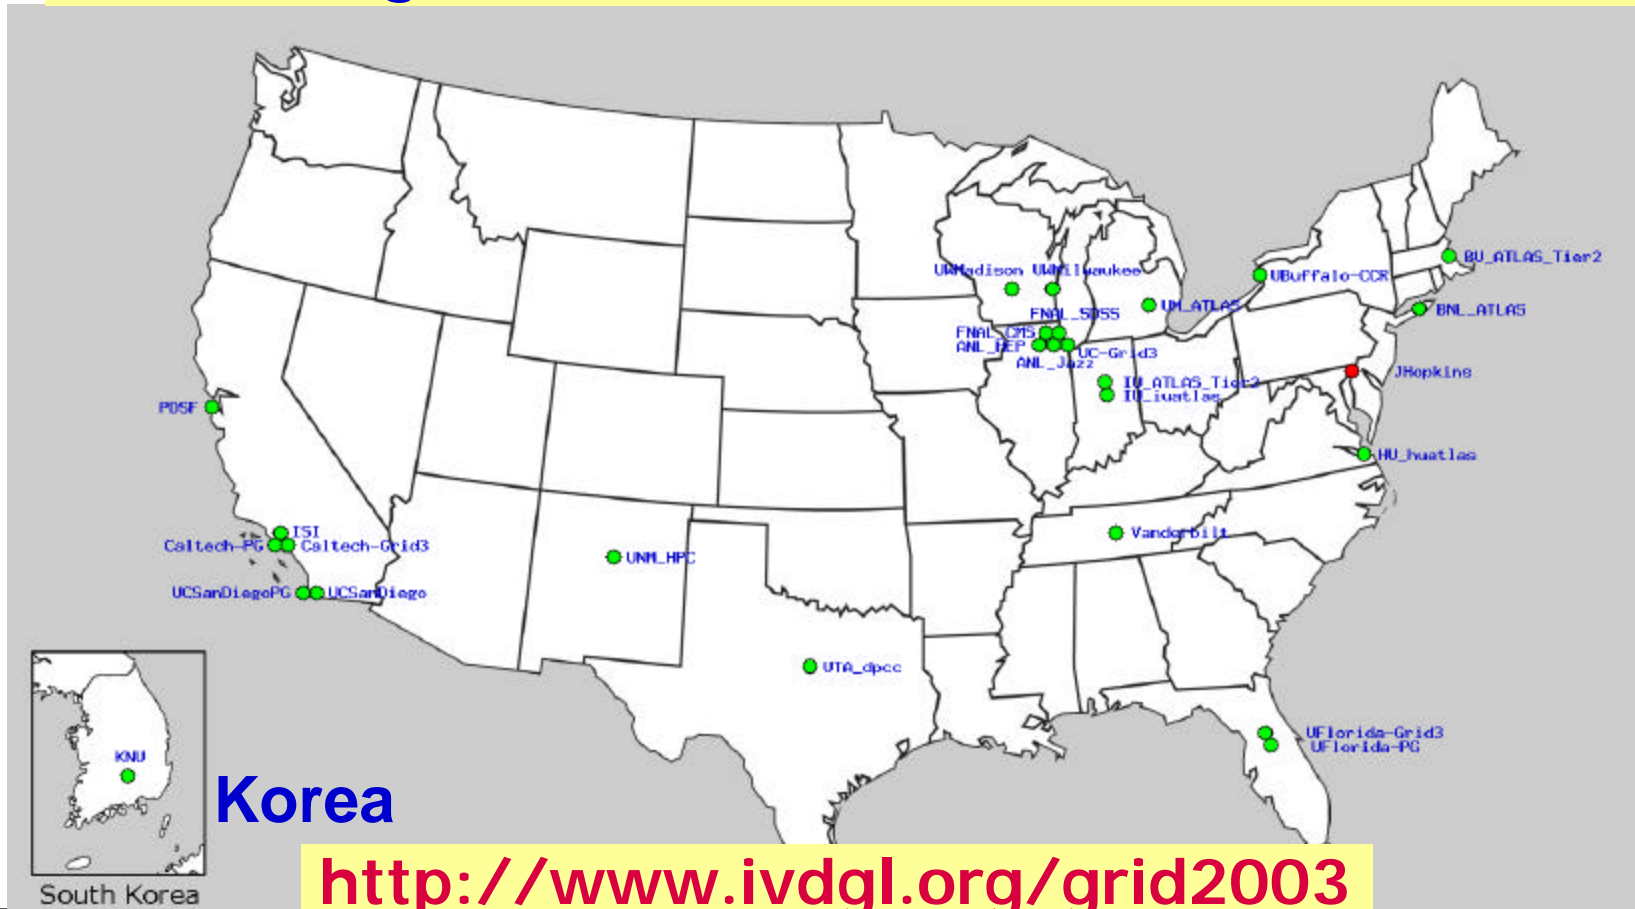
- ✍ Trillium = PPDG + GriPhyN + iVDGL
 - ✍ Large overlap in leadership, people, experiments
 - ✍ HEP members are main drivers, esp. LHC experiments
- ✍ Benefit of coordination
 - ✍ Common software base + packaging: VDT + PACMAN
 - ✍ Collaborative / joint projects: monitoring, demos, security, ...
 - ✍ Wide deployment of new technologies, e.g. Virtual Data
- ✍ Forum for US Grid projects
 - ✍ Joint strategies, meetings and work
 - ✍ Unified U.S. entity to interact with international Grid projects
- ✍ Build significant Grid infrastructure: Grid2003





Grid2003: An Operational Grid


- ~~✗~~ 28 sites (2100-2800 CPUs)
- ~~✗~~ 400-1300 concurrent jobs
- ~~✗~~ 10 applications
- ~~✗~~ Running since October 2003



<http://www.ivdgl.org/grid2003>



Grid2003 Participants





Grid2003 Applications

✂ High energy physics

- ✂ US-ATLAS analysis (DIAL),
- ✂ US-ATLAS GEANT3 simulation (GCE)
- ✂ US-CMS GEANT4 simulation (MOP)
- ✂ BTeV simulation

✂ Gravity waves

- ✂ LIGO: blind search for continuous sources

✂ Digital astronomy

- ✂ SDSS: cluster finding (maxBcg)

✂ Bioinformatics


- ✂ Bio-molecular analysis (SnB)
- ✂ Genome analysis (GADU/Gnare)

✂ CS Demonstrators

- ✂ Job Exerciser, GridFTP, NetLogger-grid2003



Grid2003: Three Months Usage






Grid2003 Success

- ✍ Much larger than originally planned
 - ✍ More sites (28), CPUs (2800), simultaneous jobs (1300)
 - ✍ More applications (10) in more diverse areas
- ✍ Able to accommodate new institutions & applications
 - ✍ U Buffalo (Biology) Nov. 2003
 - ✍ Rice U. (CMS) Feb. 2004
- ✍ Continuous operation since October 2003
 - ✍ Strong operations team (iGOC at Indiana)
 - ✍ US-CMS using it for production simulations (next slide)



Production Simulations on Grid2003

US-CMS Monte Carlo Simulation





Grid2003: A Necessary Step

✍ Learning how to operate a Grid

- ✍ Add sites, recover from errors, provide info, update, test, etc.
- ✍ Need tools, services, procedures, documentation, organization
- ✍ Need reliable, intelligent, skilled people

✍ Learning how to cope with "large" scale

- ✍ "Interesting" failure modes as scale increases
- ✍ Increasing scale must not overwhelm human resources

✍ Learning how to delegate responsibilities

- ✍ Multiple levels: Project, Virtual Org., service, site, application
- ✍ Essential for future growth

✍ **Grid2003 experience critical for building "useful" Grids**

✍ **Frank discussion in "Grid2003 Project Lessons" doc**



Grid2003 Lessons (1): Investment

✍ Building momentum

- ✍ PPDG 1999
- ✍ GriPhyN 2000
- ✍ iVDGL 2001
- ✍ Time for projects to ramp up

✍ Building collaborations

- ✍ HEP: ATLAS, CMS, Run 2, RHIC, Jlab
- ✍ Non-HEP: Computer science, LIGO, SDSS
- ✍ Time for collaboration sociology to kick in

✍ Building testbeds

- ✍ Build expertise, debug Grid software, develop Grid tools & services
- ✍ US-CMS 2002 – 2004+
- ✍ US-ATLAS 2002 – 2004+
- ✍ WorldGrid 2002 (Dec.)



Grid2003 Lessons (2): Deployment

- ✍ Building something useful draws people in
 - ✍ (Similar to a large HEP detector)
 - ✍ Cooperation, willingness to invest time, striving for excellence!
- ✍ Grid development requires significant deployments
 - ✍ Required to learn what works, what fails, what's clumsy, ...
 - ✍ Painful, but pays for itself
- ✍ Deployment provides powerful training mechanism



Grid2003 Lessons (3): Packaging

✂ Installation and configuration (VDT + Pacman)

- ✂ Simplifies installation, configuration of Grid tools + applications
- ✂ Major advances over 13 VDT releases

✂ A strategic issue and critical to us

- ✂ Provides uniformity + automation
- ✂ Lowers barriers to participation ? scaling
- ✂ Expect great improvements (Pacman 3)

✂ Automation: the next frontier

- ✂ Reduce FTE overhead, communication traffic
- ✂ Automate installation, configuration, testing, validation, updates
- ✂ Remote installation, etc.



Grid2003 and Beyond

- ✍ Further evolution of Grid3: (Grid3+, etc.)
 - ✍ Contribute resources to persistent Grid
 - ✍ Maintain development Grid, test new software releases
 - ✍ Integrate software into the persistent Grid
 - ✍ Participate in LHC data challenges
- ✍ Involvement of new sites
 - ✍ New institutions and experiments
 - ✍ New international partners (e.g., Brazil, Taiwan, ...)
- ✍ Improvements in Grid middleware and services
 - ✍ Integrating multiple VOs
 - ✍ Monitoring
 - ✍ Troubleshooting
 - ✍ Accounting
 - ✍ ...

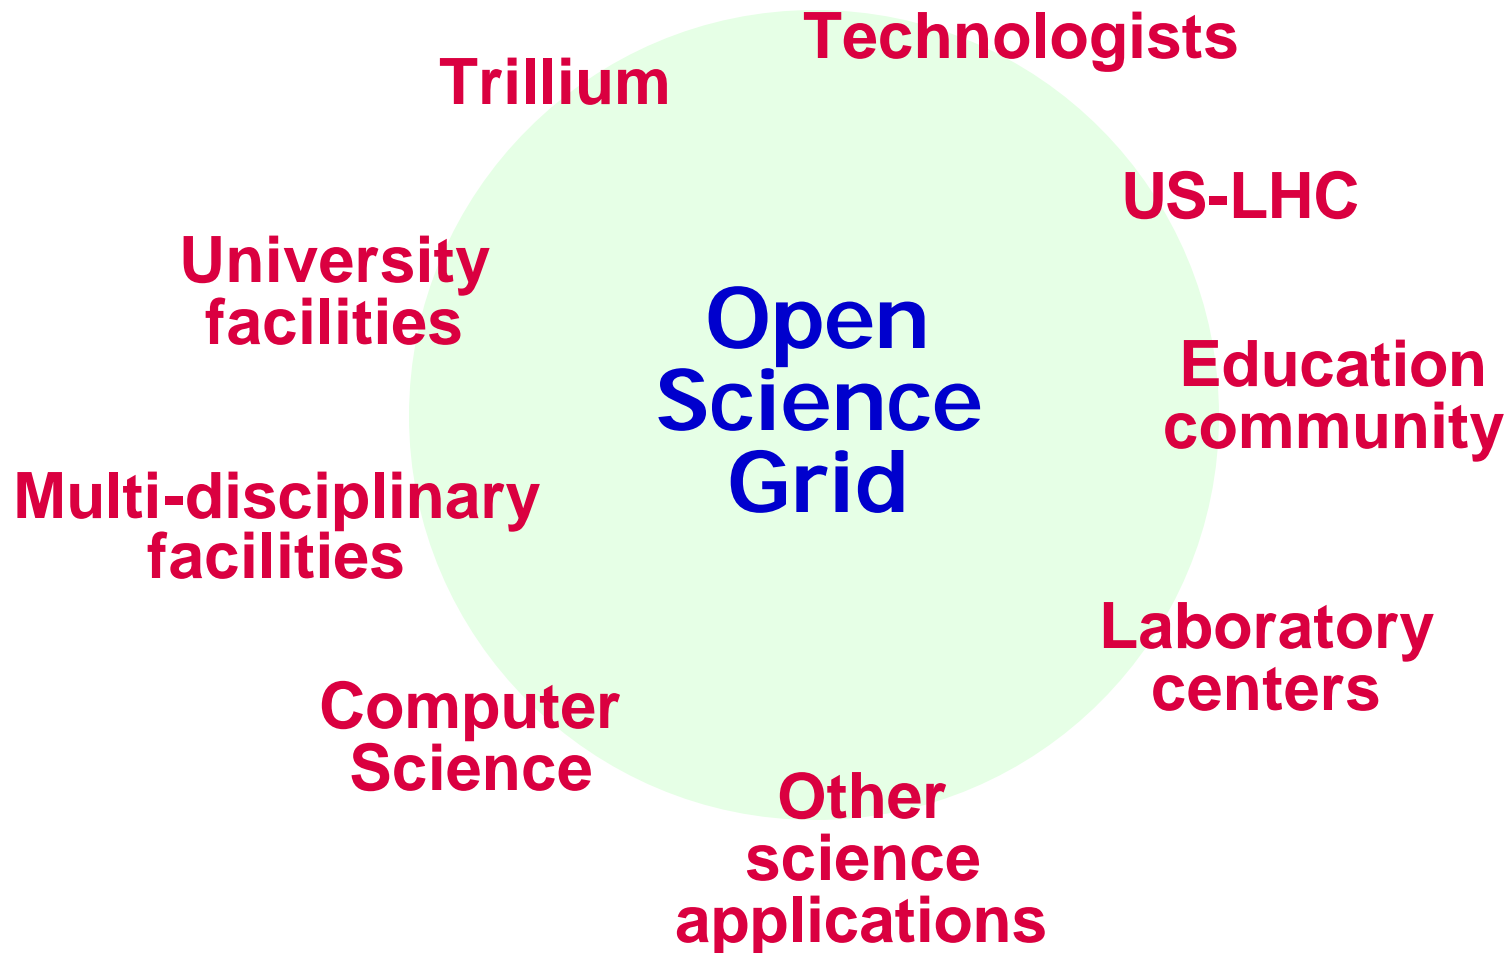


Open Science Grid

- ✍ Goal: Build an integrated Grid infrastructure
 - ✍ Support US-LHC research program, other scientific efforts
 - ✍ Resources from laboratories and universities
 - ✍ Federate with LHC Computing Grid
- ✍ Getting there: OSG-1 (Grid3+), OSG-2, ...
 - ✍ Series of releases ? increasing functionality & scale
 - ✍ Constant use of facilities for LHC production computing
- ✍ Jan. 12 meeting in Chicago
 - ✍ Public discussion, planning sessions
- ✍ Next steps
 - ✍ Creating interim Steering Committee (now)
 - ✍ White paper to be expanded into roadmap
 - ✍ Presentation to funding agencies (April/May?)



Inputs to Open Science Grid






Grids and HEP's Broad Impact

- ✍ Grids enable 21st century collaborative science
 - ✍ Linking research communities and resources for scientific discovery
 - ✍ Needed by LHC global collaborations pursuing “petascale” science
- ✍ HEP Grid projects driving important network developments
 - ✍ Recent US – Europe “land speed records”
 - ✍ ICFA-SCIC, I-HEPCCC, US-CERN link, ESNET, Internet2
- ✍ HEP has recognized leadership in Grid development
 - ✍ Many national and international initiatives
 - ✍ Partnerships with other disciplines
 - ✍ Extensive education and outreach efforts
 - ✍ Influencing funding agencies (NSF, DOE, EU)
- ✍ Grid2003 and its successors will provide a scalable computing infrastructure for scientific research



Grid References

- ✍ Grid2003
 - ✍ www.ivdgl.org/grid2003
- ✍ Globus
 - ✍ www.globus.org
- ✍ PPDG
 - ✍ www.ppdg.net
- ✍ GriPhyN
 - ✍ www.griphyn.org
- ✍ iVDGL
 - ✍ www.ivdgl.org
- ✍ LCG
 - ✍ www.cern.ch/lcg
- ✍ EU DataGrid
 - ✍ www.eu-datagrid.org
- ✍ EGEE
 - ✍ egee-ei.web.cern.ch





Extra Slides




Grid Challenges

- ✍ Operate a fundamentally complex entity
 - ✍ Geographically distributed resources
 - ✍ Each resource under different administrative control
 - ✍ Many failure modes
- ✍ Manage workflow of 1000s of jobs across Grid
 - ✍ Balance policy vs. instantaneous capability to complete tasks
 - ✍ Balance effective resource use vs. fast turnaround for priority jobs
 - ✍ Match resource usage to policy over the long term
- ✍ Maintain a global view of resources and system state
 - ✍ Coherent end-to-end system monitoring
- ✍ Build managed system & integrated user environment
 - ✍ Integrating computing, storage, networks
 - ✍ Providing transparent usage of experiment resources




LHC: Key Driver for Data Grids

- ✍ Complexity: Millions of individual detector channels
- ✍ Scale: PetaOps (CPU), 100s of Petabytes (Data)
- ✍ Distribution: Global distribution of people & resources






Most IT Resources Outside CERN






LHC Global Collaborations

ATLAS



CMS



US has ~20-25% of physicists

**1849 Physicists and Engineers
34 Countries
147 Institutions**

LHC Data and CPU Requirements

ATLAS

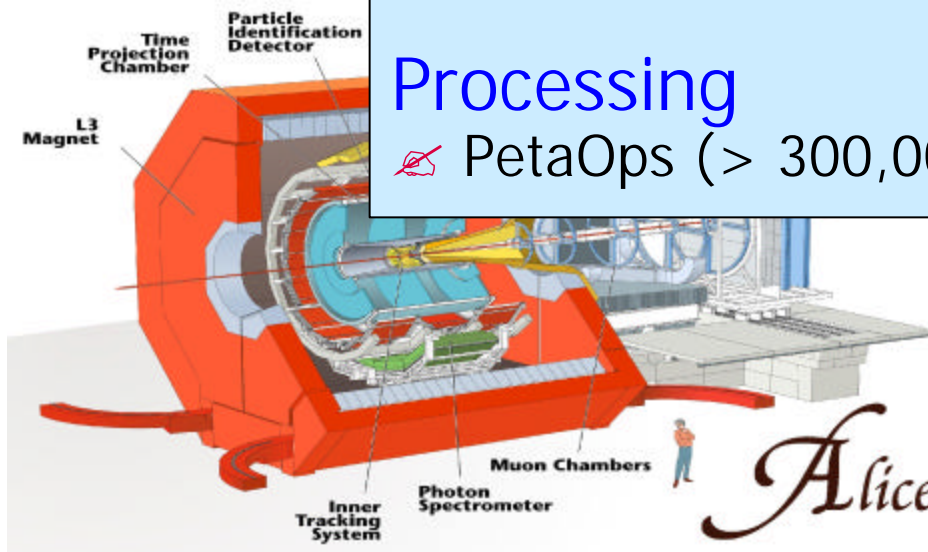
CMS

Storage


- ~~✗~~ Raw recording rate 0.1 – 1 GB/s
- ~~✗~~ Monte Carlo data
- ~~✗~~ 100 PB total by ~2012

Processing

- ~~✗~~ PetaOps (> 300,000 3 GHz PCs)



ALICE





Driver for Transatlantic Networks (Gb/s)

	2001	2002	2003	2004	2005	2006
<i>CMS</i>	0.1	0.2	0.3	0.6	0.8	2.5
<i>ATLAS</i>	0.05	0.1	0.3	0.6	0.8	2.5
<i>BaBar</i>	0.3	0.6	1.1	1.6	2.3	3.0
<i>CDF</i>	0.1	0.3	0.4	2.0	3.0	6.0
<i>D0</i>	0.4	1.6	2.4	3.2	6.4	8.0
<i>BTeV</i>	0.02	0.04	0.1	0.2	0.3	0.5
<i>DESY</i>	0.1	0.18	0.21	0.24	0.27	0.3
<i>CERN</i>	0.31	0.62	2.5	5.0	10	20

2001 estimates, now seen as conservative!



HEP Bandwidth Roadmap (Gb/s)

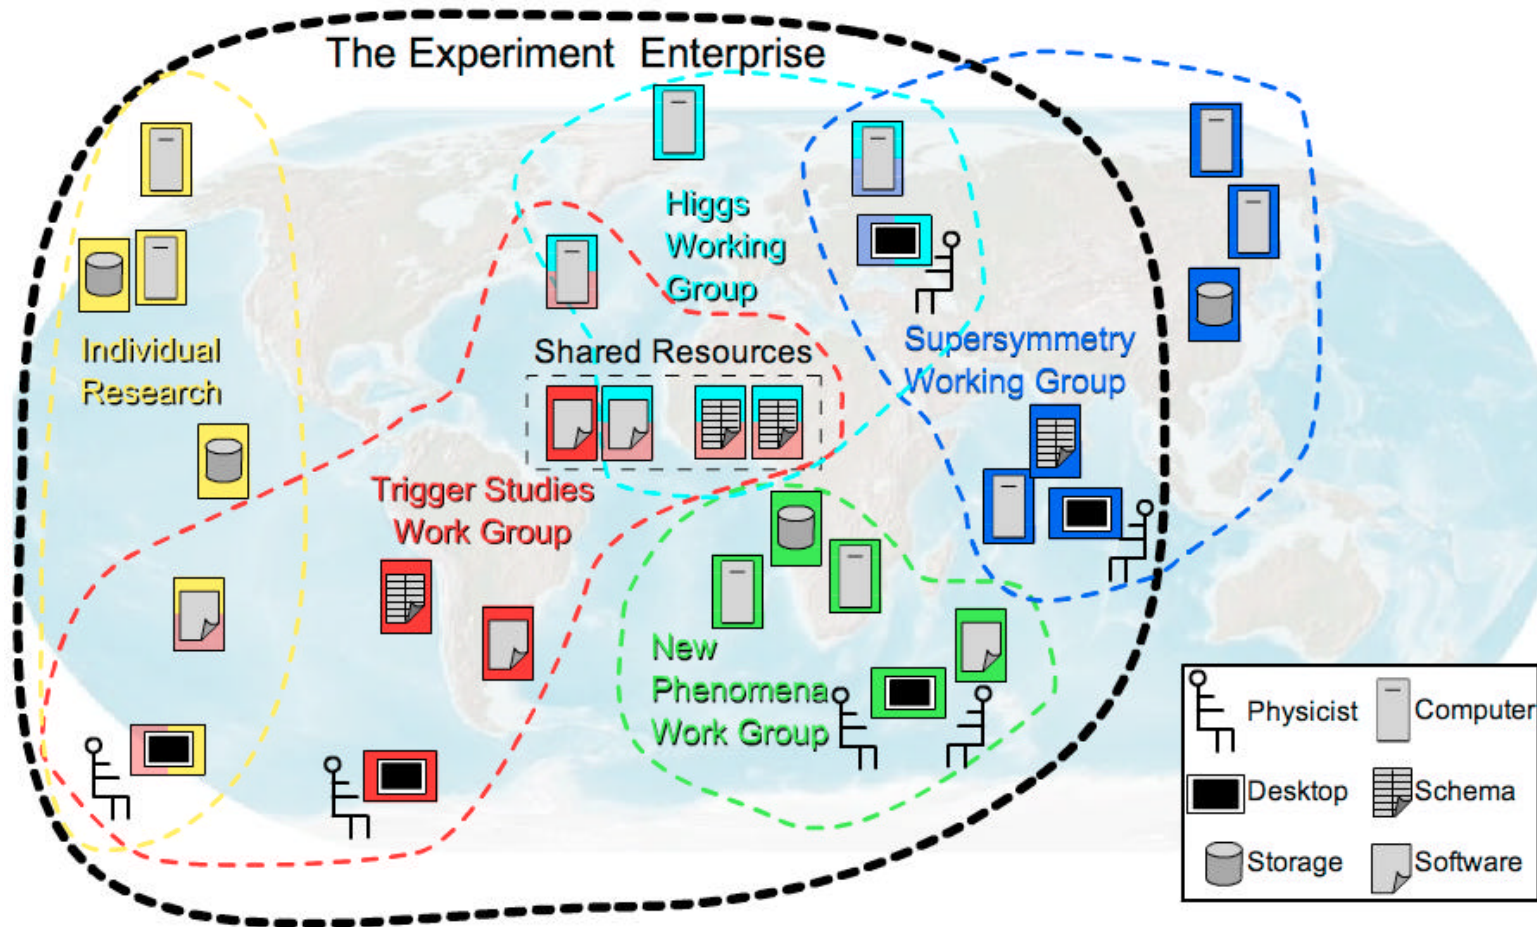
<i>Year</i>	<i>Production</i>	<i>Experimental</i>	<i>Remarks</i>
2001	0.155	0.62 - 2.5	SONET/SDH
2002	0.622	2.5	SONET/SDH DWDM; GigE Integ.
2003	2.5	10	DWDM; 1+10 GigE Integration
2005	10	2-4 ? 10	? Switch; ? Provisioning
2007	2-4 ? 10	1 ? 40	1 st Gen. ? Grids
2009	1-2 ? 40	~5 ? 40	40 Gb/s; ? Switching
2011	~5 ? 40	~25 ? 40	Terabit Networks
2013	~Terabit	~MultiTbps	~Fill One Fiber

- ✍ Continuing trend of ? 1000 bandwidth growth per decade?
- ✍ HEP plays leading role in future networks



Analysis by Globally Distributed Teams

- ~~Non-hierarchical: Chaotic analyses + productions~~
- ~~Superimpose significant random data flows~~





Grids: Enhancing Research & Learning

✍ Fundamentally alters conduct of scientific research



✍ "Lab-centric": Activities center around large facility

✍ "Team-centric": Resources shared by distributed teams

✍ "Knowledge-centric": Knowledge generated/used by a community

✍ Strengthens role of universities in research

✍ Couples universities to data intensive science

✍ Couples universities to national & international labs

✍ Brings front-line research and resources to students

✍ Exploits intellectual resources of formerly isolated schools

✍ Opens new opportunities for minority and women researchers

✍ Builds partnerships to drive advances in IT/science/eng

✍ HEP ? Physics, astronomy, biology, CS, etc.

✍ "Domain" sciences ? Computer Science

✍ Universities ? Laboratories

✍ Scientists ? Students

✍ Research Community ? IT industry





UNIVERSITY OF
FLORIDA



Universidade do Estado
do Rio de Janeiro

CHEPREO: Center for High Energy Physics Research and Educational Outreach



- ✍ Physics Learning Center (Miami)
- ✍ CMS research
- ✍ iVDGL Grid activities
- ✍ AMPATH network (S. America)

Funded September 2003
\$4M (MPS, CISE, EHR, ENG)



U.S. "Trillium" Physics Grid Projects

✍ Trillium = PPDG + GriPhyN + iVDGL

✍ PPDG: \$10M (DOE) (1999 – 2004+)

✍ GriPhyN: \$12M (NSF) (2000 – 2005)

✍ iVDGL: \$14M (NSF) (2001 – 2006)



✍ Basic composition (~150 people)

✍ PPDG: 10 universities, SDSC, 5 labs

✍ GriPhyN: 12 universities, SDSC, 4 labs

✍ iVDGL: 18 universities, SDSC, 4 labs, foreign partners

✍ Expts: BaBar, D0, STAR, Jlab, CMS, ATLAS, LIGO, SDSS/NVO

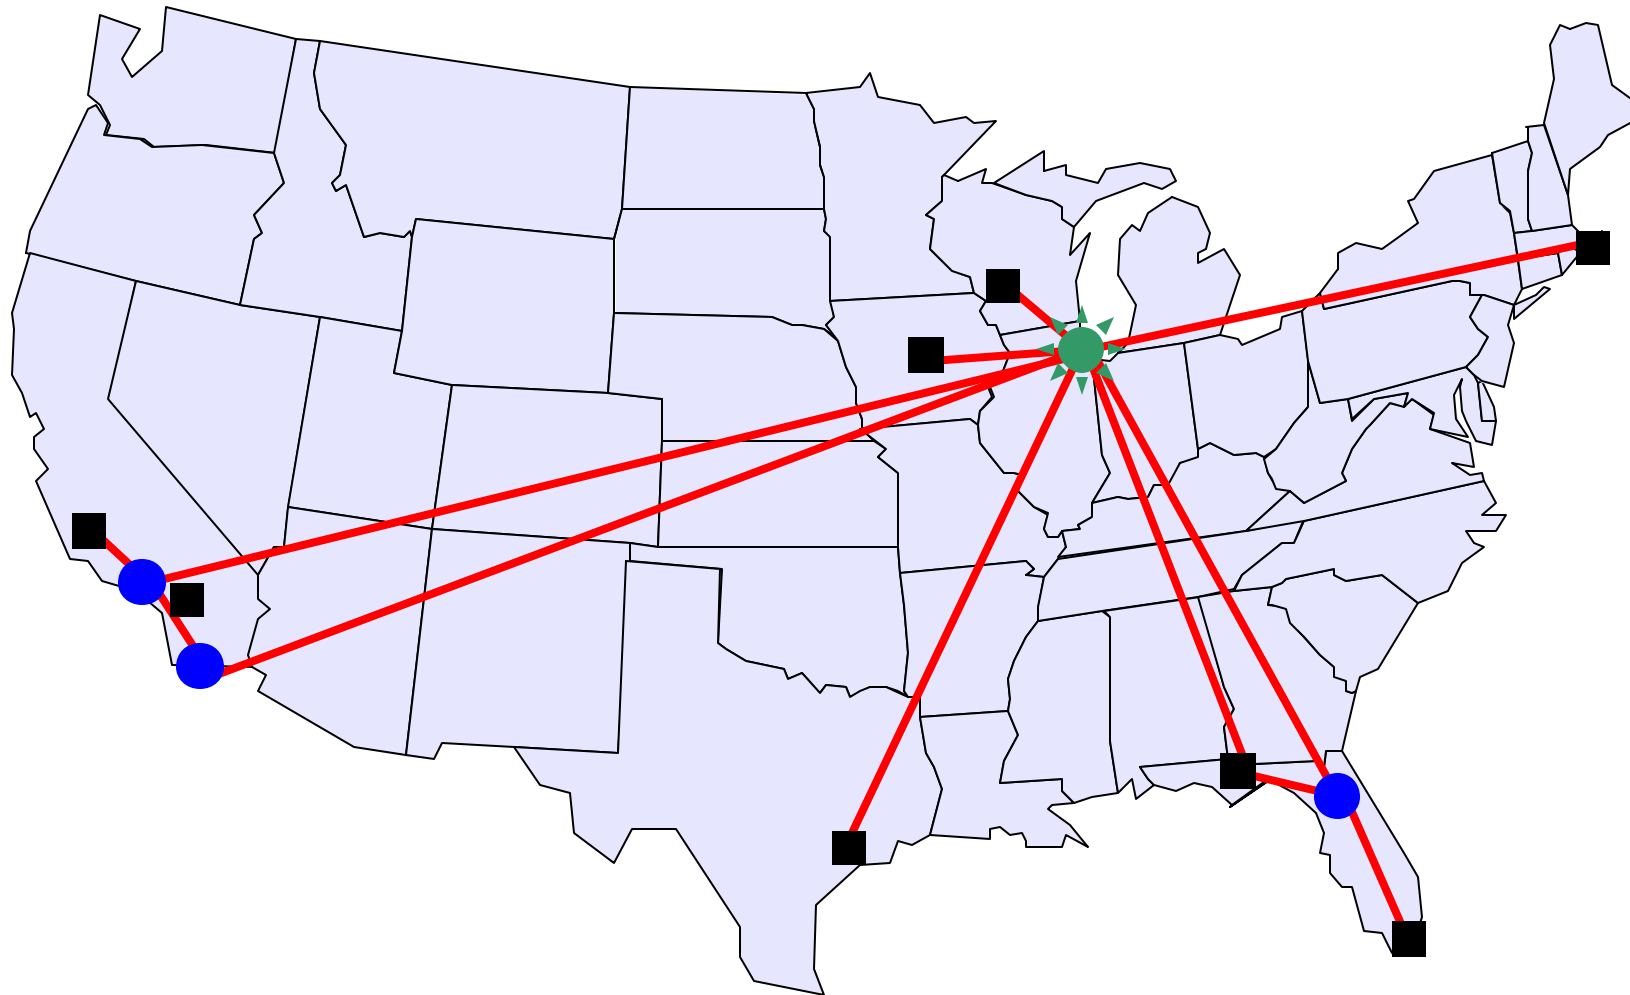
✍ Grid research + infrastructure + deployment

✍ GriPhyN: CS research, Virtual Data Toolkit (VDT) development

✍ PPDG: "End to end" Grid, services, monitoring, analysis

✍ iVDGL: Grid laboratory deployment, university Tier-2 resources

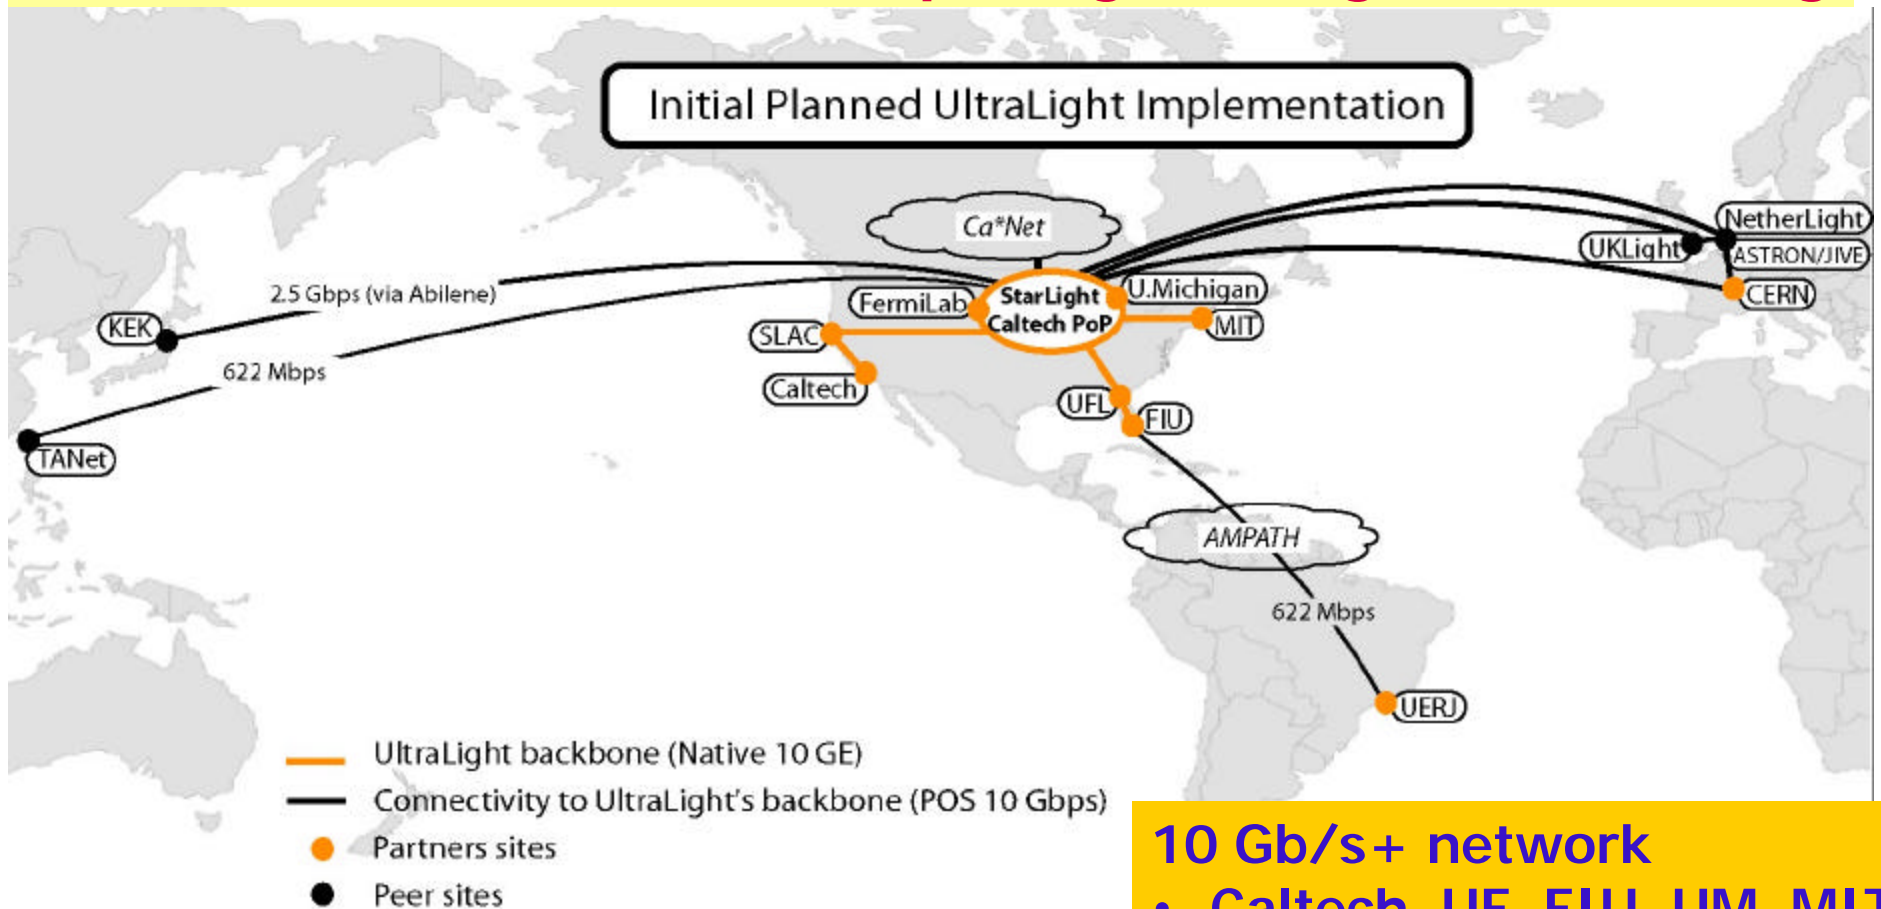
✍ HENP experiments provide frontier challenges





UltraLight

Unified Infrastructure: Computing, Storage, Networking



10 Gb/s+ network

- Caltech, UF, FIU, UM, MIT
- SLAC, FNAL, BNL
- Int'l partners
- Cisco, Level(3), Internet2



Grid2003: A Path to Open Science Grid?

✍ Yes

✍ Grid2003 proved that a prototype production Grid could be built

✍ ... but much more is needed

✍ Security, user account management

✍ Multiple Virtual Organizations

✍ Storage and cluster management

✍ Ubiquitous monitoring, accounting

✍ Scheduling

✍ Integration of databases, optical networks

✍ Heterogeneity (IA64, other Linux flavors)

✍ More and different kinds of applications

✍ Need manpower + computing resources



LCG: LHC Computing Grid Project

- ✍ A persistent Grid infrastructure for LHC experiments
 - ✍ Matched to decades long research program of LHC
- ✍ Prepare & deploy computing environment for LHC expts
 - ✍ Common applications, tools, frameworks and environments
 - ✍ *Deployment only*: no middleware development
- ✍ Move from testbed systems to real production services
 - ✍ Operated and supported 24x7 globally
 - ✍ Computing fabrics run as production physics services
 - ✍ A robust, stable, predictable, supportable infrastructure





LCG Activities

✍ Relies on strong leverage

- ✍ CERN IT
- ✍ LHC experiments: ALICE, ATLAS, CMS, LHCb
- ✍ Grid projects: Trillium, EU DataGrid, EGEE

✍ Close relationship with EGEE project

- ✍ EGEE will create EU production Grid (funds 70 "partners")
- ✍ Middleware activities done in common (same manager)
- ✍ LCG deployment in common (LCG Grid is EGEE prototype)

✍ LCG-1 deployed (Sep. 2003)

- ✍ Tier-1 laboratories (including FNAL and BNL)

✍ LCG-2 in process of being deployed



Organisation Européenne pour la Recherche Nucléaire
European Organization for Nuclear Research

Sep. 29, 2003 announcement

PR13.0
29.09.200



Renovation of the Computer Centre
at CERN at this moment which
"looks like a grid"...

LHC Computing Grid Goes Online



The world's particle physics community today announced the launch of the first phase of the [LHC computing Grid \(LCG\)](#). The LCG is designed to handle the unprecedented quantities of data that will be produced by experiments at CERN's [Large Hadron Collider \(LHC\)](#) from 2007 onwards.

"The LCG will provide a vital test-bed for the new Grid computing technologies that are set to revolutionise the way scientists use the world's computing resources in areas ranging from fundamental research to medical diagnosis," said [Les Robertson](#), CERN's LCG project manager.

The computational requirements of the experiments that will operate at the LHC are enormous. Some 12-14 petabytes of data will be generated each year, the equivalent of more than 20 million CDs. Analysing this data will require the equivalent of 70,000 of today's fastest PC computers. The LCG will meet these needs by deploying a worldwide computational Grid, integrating the resources of scientific computing centres spread across Europe, America and Asia into a global virtual computing service.

The first phase of the project, LCG-1, will operate a series of prototype services, gradually increasing in scale and complexity as its builders develop an understanding of the functional and operational complexities involved in building a Grid of such unprecedented scale. LCG-1 uses so-called 'middleware' developed mainly by the [European Data Grid](#) project in Europe and the Globus, Condor and related projects contributing to the Virtual Data Toolkit in the US. It allows physicists to access worldwide distributed computing resources from their desktops as if they were local.



Current LCG-1 Sites





Grid2003 and Beyond (2)

- ✍ Coordination with LHC Computing Grid
 - ✍ Federating Grid2003 resources with LCG
 - ✍ Participate in global “data challenge” exercises
- ✍ Involvement of other disciplines
 - ✍ CS, LIGO, Astronomy, Biology, ...
- ✍ Coordinated education/outreach
 - ✍ QuarkNet, GriPhyN, iVDGL, PPDG, CMS, ATLAS
 - ✍ CHEPREO center at Florida International U
 - ✍ Digital Divide efforts (Feb. 15-20 Rio workshop)